



t-DGR: A Trajectory-Based Deep Generative Replay Method for Continual Learning in Decision Making

William Yue¹, Bo Liu¹, Peter Stone^{1,2}

The University of Texas at Austin¹, Sony AI²




Continual Learning and Deep Generative Replay



Deep generative replay in machine learning involves training a neural network to generate samples from past experiences, mitigating catastrophic forgetting in continual learning,


Continual World Benchmark Results



The Continual World CW20 benchmark consists of 10 tasks, each repeated twice, where a Sawyer arm manipulates one or two objects in the MuJoCo physics simulator.

Method	Success Rate \uparrow	FT \uparrow	Forgetting \downarrow
Finetune	14.2 \pm 4.0	-0.5 \pm 3.0	82.2 \pm 5.6
Multitask	97.0 \pm 1.0	N/A	N/A
oEWC	19.4 \pm 5.3	-2.8 \pm 4.1	75.2 \pm 7.5
PackNet	74.1 \pm 4.1	-20.4 \pm 3.4	-0.2 \pm 0.9
DGR	74.1 \pm 4.1	18.9 \pm 2.9	23.3 \pm 3.3
CRIL	50.8 \pm 4.4	4.4 \pm 4.9	46.1 \pm 5.4
t-DGR	83.9 \pm 3.0	30.6 \pm 4.5	14.6 \pm 2.9

Ground Truth





CRIL



DGR




t-DGR (ours)



We evaluate three methods for replicating an agent's movement in a 2D plane. CRIL deviates at a sharp turn and **never recovers**, while DGR produces a **fragmented path**. Our proposed method, t-DGR, overcomes these issues by sampling state observations based on the **trajectory timestep**.

Current Pseudo-Rehearsal Methods



Generates video frames from the distribution observed during training, resulting in a **bag of unordered video frames**.

Generates the next video frame by **conditioning on the preceding video frame**.

Generates a video frame by **conditioning on the trajectory timestep**.